

THIS OPINION WAS NOT WRITTEN FOR PUBLICATION

The opinion in support of the decision being entered today (1) was not written for publication in a law journal and (2) is not binding precedent of the Board.

Paper No. 23

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE BOARD OF PATENT APPEALS
AND INTERFERENCES

Ex parte ODED COHN, ELIEZER DEKEL,
and MICHAEL RODEH

MAILED

JAN 28 2000

Appeal No. 1997-4123
Application No. 08/364,334¹

PAT. & T.M. OFFICE
BOARD OF PATENT APPEALS
AND INTERFERENCES

ON BRIEF

Before THOMAS, HAIRSTON, and BARRY, Administrative Patent Judges.
BARRY, Administrative Patent Judge.

DECISION ON APPEAL

This is a decision on the appeal under 35 U.S.C. § 134 from the final rejection of claims 1-3 and 10-18. We reverse.

¹ The application was filed on December 27, 1994.

BACKGROUND

The invention at issue in this appeal reduces the power consumed by a battery-powered computer that employs a cache and disk drive. In a typical battery-powered computer, the disk drive of the computer is set to an idle mode when the disk has not been accessed for a certain time. Upon receipt of a read or write request to the disk, the disk is powered-up to an active mode. Powering-up the disk, however, consumes much power.

The use of a cache to supplement a disk drive in a computer is well known. To maintain consistency between the cache and disk, data written into the cache must be written back to the disk. Prior art computers perform such write-back in the background, i.e., independent of requests to the disk. Such background operations independently power-up the disk, however, thereby consuming additional power.

The invention reduces the number of times a disk is powered-up, thereby reducing power consumption. Specifically, data are written-back to the disk only when the disk has been powered-up to service a read or write request.

Claim 18, which is representative for our purposes, follows:

18. A method for operating data storage apparatus having a storage element including at least one moveable part and a solid-state cache memory in which the storage element is in a non-fully operational mode when data access has not occurred during a predetermined time period, the method comprising the steps of:

accessing data on the storage element in response to a read or write request which cannot be satisfied via access to the cache memory;

accessing data stored within the cache memory if a read or write request from the computer can be satisfied via access to the cache memory;

designated and selected data within said cache memory as new data in response to a write request from the computer which updates data within the cache memory; and

transferring new data from the cache memory to the storage element to maintain consistency of data therebetween, said data transfers between the cache memory and the storage element being performed when the storage element is fully operational as a result of a read or write request has given rise to an access to the storage element.

The references relied on in rejecting the claims follow:

Hanson et al. 4,433,374 Feb. 21, 1984
(Hanson)

Noya et al. 5,420,983 May 30, 1995
(Noya) (filed Aug. 12, 1992)

Yamazaki² 4-205852 Nov. 30, 1990.

² A copy of the translation prepared by the U.S. Patent and (continued...)

Claims 1, 2, 10-15, and 18 stand rejected under 35 U.S.C. § 103 as obvious over Yamazaki. Claim 3 stands rejected under 35 U.S.C. § 103 as obvious over Yamazaki in view of Hanson. Claims 16 and 17 stand rejected under 35 U.S.C. § 103 as obvious over Yamazaki in view of Noya. Rather than repeat the arguments of the appellants or examiner in toto, we refer the reader to the brief and answer for the respective details thereof.

OPINION

In reaching our decision in this appeal, we considered the subject matter on appeal and the rejections advanced by the examiner. Furthermore, we duly considered the arguments and evidence of the appellants and examiner. After considering the totality of the record, we are persuaded that the examiner erred in rejecting claims 1-3 and 10-18. Accordingly, we reverse.

We begin by noting the following principles from In re Rijckaert, 9 F.3d 1531, 1532, 28 USPQ2d 1955, 1956 (Fed. Cir. 1993).

²(...continued)
Trademark Office is included and relied upon for this decision. We will refer to the translation by page number in this opinion.

In rejecting claims under 35 U.S.C. Section 103, the examiner bears the initial burden of presenting a prima facie case of obviousness. In re Oetiker, 977 F.2d 1443, 1445, 24 USPQ2d 1443, 1444 (Fed. Cir. 1992). Only if that burden is met, does the burden of coming forward with evidence or argument shift to the applicant. Id. "A prima facie case of obviousness is established when the teachings from the prior art itself would appear to have suggested the claimed subject matter to a person of ordinary skill in the art." In re Bell, 991 F.2d 781, 782, 26 USPQ2d 1529, 1531 (Fed. Cir. 1993) (quoting In re Rinehart, 531 F.2d 1048, 1051, 189 USPQ 143, 147 (CCPA 1976)). If the examiner fails to establish a prima facie case, the rejection is improper and will be overturned. In re Fine, 837 F.2d 1071, 1074, 5 USPQ2d 1596, 1598 (Fed. Cir. 1988).

With these in mind, we analyze the appellants' argument.

The appellants argue, "Yamazaki is absolutely silent on the subject of modified data within the cache and how that data is transferred to the disk storage element" (Appeal Br. at 7.) The examiner makes the following reply.

[A]n artisan would be aware of the competing interests of maximizing data safety and minimizing power conservation. Data safety is enhance[d] by quickly flushing dirty entries, while power consumption is maximized during the spin up process and minimized in the idle state. Consequently, the Examiner believes that an artisan would find it obvious to destage data from cache to disk only when the storage element has been brought to operating speed as a result of another read/write transaction, since this would serve the competing interests of maximizing data safety and minimizing power consumption. (Examiner's Answer at 6.)

We agree with the appellants.

The examiner errs in interpreting the claims. He characterizes the invention as "a cache replacement mechanism which flushes dirty entries when another disk transaction requires access to the disk itself" (Final Rejection at 3.) In contrast, claims 1-3 and 10-17 each specifies in pertinent part the following limitations:

a storage element ... wherein the storage element has an operating, and a non fully operating mode when data access has not occurred for a predetermined time period;

...

means for designating selected data within said cache memory as new data in response to a write request from the computer which updates data within the cache memory; and

a cache replacement mechanism for transferring new data from the cache memory to the storage element to maintain consistency of data between data stored in the cache memory and data stored in the storage element, wherein the cache replacement mechanism performs data transfers between the cache memory and the storage element when the storage element is at operating speed as a result of a read or write request which requires an access to the storage element.

Similarly, claim 18 specifies in pertinent part the following limitations:

18. A method for operating data storage apparatus having a storage element ... in which the storage

element is in a non-fully operational mode when data access has not occurred during a predetermined time period, the method comprising the steps of:

...

designated and selected data within said cache memory as new data in response to a write request from the computer which updates data within the cache memory; and

transferring new data from the cache memory to the storage element to maintain consistency of data therebetween ... when the storage element is fully operational as a result of a read or write request has given rise to an access to the storage element.

In summary, the limitations recite writing-back data from a cache to a disk only when the disk has been powered-up to service a read or write request.

"Obviousness may not be established using hindsight or in view of the teachings or suggestions of the inventor."

Para-Ordnance Mfg. v. SGS Importers Int'l, 73 F.3d 1085, 1087, 37 USPQ2d 1237, 1239 (Fed. Cir. 1995) (citing W.L. Gore & Assocs., Inc. v. Garlock, Inc., 721 F.2d 1540, 1551, 1553, 220 USPQ 303, 311, 312-13 (Fed. Cir. 1983), cert. denied, 469 U.S. 851 (1984)). "A rejection based on section 103 clearly must rest on a factual basis" In re Warner, 379 F.2d 1011, 1017, 154 USPQ 173, 178 (CCPA 1967). "The Patent Office has the initial duty of supplying the factual basis for its rejection. It may

not ... resort to speculation, unfounded assumptions[,] or hindsight reconstruction to supply deficiencies in its factual basis." Id., 154 USPQ at 178.

The examiner fails to produce any evidence to support his allegation that Yamazaki, which is the sole reference relied on in rejecting the independent claims, "inherently teaches a means for designating selected data within the cache as new data in response to a write from the computer which updates data within the cache" (Final Rejection at 2-3.) The reference shows employing a cache 204 to store data being read from a magnetic disk, i.e., employing the cache as a read cache. Translation, Fig. 1. In contrast, Yamazaki does not mention employing the cache 204 to store data being written to the magnetic disk, i.e., employing the cache as a write cache. By teaching a read cache 40 that is separate from a write cache 50, moreover, Noya evidences that a cache does not necessarily operate as both a read cache and a write cache. Col. 5, ll. 2-4. In view of Yamazaki's omission and Noya's teaching, the examiner's allegation amounts to speculation or an unfounded assumption.

The examiner also fails to produce any evidence to support his conclusion that "an artisan would find it obvious to destage data from cache to disk only when the storage element has been brought to operating speed as a result of another read/write transaction" (Examiner's Answer at 6.) Yamazaki fails to teach writing-back data from a cache to a disk at all. Although Hanson and Noya each teaches writing-back data from a cache to a disk, neither reference teaches waiting until the disk has been powered-up to service a read or write request to write-back the cached data. In view of these failures, the examiner's conclusion amounts to impermissible reliance on the appellants' teachings or suggestions.

For the foregoing reasons, we are not persuaded that the prior art would have suggested writing-back data from a cache to a disk only when the disk has been powered-up to service a read or write request as claimed. The examiner has not established a prima facie case of obviousness. Therefore, we reverse the rejection of claims 1-3 and 10-18 under 35 U.S.C. § 103.

CONCLUSION

To summarize, the examiner's rejection of claims 1-3 and 10-18 under 35 U.S.C. § 103 is reversed.

REVERSED

JAMES D. THOMAS)
Administrative Patent Judge)

KENNETH W. HAIRSTON)
Administrative Patent Judge)

LANCE LEONARD BARRY)
Administrative Patent Judge)

)
BOARD OF PATENT
APPEALS
AND
INTERFERENCES

Appeal No. 1997-4123
Application No. 08/364, 334

Page 11

ANDREW J. DILLON
FELSMAN, BRADLEY, GUNTER
& DILLON, LLP
SUITE 350 ARBORETUM POINT
9505 ARBORETUM BOULEVARD
AUSTIN, TX 78759

TRANSLATION FROM JAPANESE

- (19) JAPANESE PATENT OFFICE (JP)
(11) Japanese Laid-Open Patent Application (Kokai) No. 4-205852
(12) Official Gazette for Laid-Open Patent Applications (A)
- (51) Int. Cl.⁵: Classification Symbols: Internal Office Registration Nos.:

G 11 B	19/00	H	7627-5D
G 06 F	1/32		7832-5B
G 06 F	1/00	332	E
G 06 F	12/08	320	7232-5B

(43) Laying-Open Date: July 28, 1992
Request for Examination: Not yet submitted
Number of Claims: 4
(Total of 5 pages [in original])

- (54) Title of the Invention: Disk Cache Device
- (21) Application No. 2-328844
(22) Filing Date: November 30, 1990
(72) Inventor: Yukie Yamazaki
(72) Inventor: Yuji Shimada
(72) Inventor: Yasuo Hatchi
(72) Inventor: Yoshito Suzuki
(71) Applicant: Hitachi, Ltd.
(71) Applicant: Hitachi Micro Software Systems K.K.
(74) Agent: Katsuo Ogawa, Patent Attorney (and one other)

SPECIFICATION

1. Title of the Invention

Disk Cache Device

2. Claims

1. A disk cache device, comprising a magnetic disk, a cache memory, and a controller for controlling said magnetic disk and cache memory, wherein said disk cache device is characterized by the fact that said controller determines whether or not said cache memory contains data for which an access request has been received from a host computer, actuates the spindle motor of said magnetic disk only when [said data] are not present in said cache memory, and again stops said spindle motor after access to said magnetic disk has been completed.
2. A disk cache device as defined in Claim 1, characterized by the fact that said controller stops the operation of the spindle motor of said magnetic disk when said cache memory contains data for which an access request has been received from said host computer.
3. A disk cache device as defined in Claim 1 or 2, wherein said spindle motor of said magnetic disk is not switched between operation and stoppage when the hit rate is lower than a separately determined value.
4. A disk cache device as defined in Claim 2, wherein said controller actuates said spindle motor of said magnetic disk when data for which an access request has been received from said host computer are not present in said cache memory.

3. Detailed Description of the Invention

Field of Industrial Utilization

The present invention relates to a disk device having a cache memory.

As described in the conventional examples cited in the DRPO 40C product specifications for hard disks manufactured by Alps Electric, commands for placing a magnetic disk in a low power consumption mode (and thus to lower the power consumption of the magnetic disk device) have in the past been issued by a host computer. An alternative has been to allow a magnetic disk itself to automatically switch to a low power consumption mode when no commands have been received by the magnetic disk over a predetermined period of time.

Problems Which the Invention Is Intended to Solve

With the aforementioned prior art, no consideration was given to the power consumption of a magnetic disk having a disk cache, and no provisions were made for reducing the power consumption of such magnetic disks. For a magnetic disk having a disk cache, however, a request for accessing the disk from the host computer (even when such a request is made) does not require accessing the disk while the corresponding data are held in the cache memory, that is, while a hit is made, and the power consumed by the disk in this period is wasted. It is therefore desirable to place the disk in a low power consumption mode when the requested data have been hit, even if an access request has been received from the host.

It is, however, impossible for the host computer to determine whether the requested data have been hit or missed, so the host computer cannot place the disk in a low power consumption mode when the requested data are hit, and electric power is wasted.

An object of the present invention is to lower the power consumption of a magnetic disk device having a disk cache.

Means Used to Solve the Above-Mentioned Problems

Aimed at attaining the stated objective, the present invention involves actuating a magnetic disk only when the cache memory does not contain the data requested by the host (when a miss occurs) and the magnetic disk is accessed, and placing the magnetic disk in a low power consumption mode under all other conditions. Specifically, the magnetic disk is

switched between different modes of operation by a magnetic disk controller independently of the host computer.

As used herein, the term "low power consumption mode of a magnetic disk" refers either to a mode in which power supply to ENDEC, an R/W circuit, or another signal processing circuit has been cut and the spindle motor of the disk stopped, or to a mode in which power supply to a signal processing circuit has been cut, but the spindle motor continues to operate.

Effect of the Invention

Because the magnetic disk controller controls the cache memory, it is clear whether or not the data requested by the host have been hit or missed. The fact that the magnetic disk controller switches the magnetic disk between an operating mode and a low power consumption mode makes it possible to place the magnetic disk in the low power consumption mode when the data have been hit.

Practical Examples

A practical example of the present invention will now be described with reference to figures.

Figure 2 is a block diagram of the hardware pertaining to a practical example of the present invention.

In Figure 2, 201 is a host computer, 202 a host interface controller, 203 a disk/cache memory controller, and 204 a cache memory.

An encoder/decoder data separator (205), a waveform shaping circuit (206), and an R/W circuit (207) are components of a signal processing circuit; the disk/cache memory controller 203 can control the feeding and cutting of power via power save control 2.

210 is a spindle motor drive. The disk/cache memory controller 203 can control the manner in which a spindle motor 211 is switched on and off via the spindle motor drive 210 by means of power save control 1.

A processing run performed in accordance with the present invention will now be described with reference to Figure 1. The processing starts with step 11. In step 12, the disk/cache memory controller 203 switches off the spindle motor 211 using power save control 2. In conjunction with this, the power supply of the signal processing circuit is switched off using power save control 1.

Next, in step 13, the disk/cache memory controller 203 is held in a standby mode of low power consumption until a request is received from the host. As referred to herein, the low power consumption mode may, for example, involve stopping the input of clock [pulses] to the disk/cache memory controller, reducing the loss of electric power.

In step 14, in which a request is received from the host, the disk/cache memory controller 203 cancels the low power consumption mode.

Next, in step 15, the disk/cache memory controller 203 determines whether or not the data requested by the host have been hit or missed. In the case of a hit, data are exchanged between the cache memory and the host during step 16. If a miss occurred during step 15, the spindle motor 211 of the magnetic disk is switched on during step 17, and power is fed to the signal processing circuit of the magnetic disk. In step 18, data are exchanged between the host and the magnetic disk via the cache memory 204, and in step 19 the spindle motor 211 of the magnetic disk is switched off again, and power is cut off from the signal processing circuit of the magnetic disk.

Another practical example will now be described with reference to Figure 3.

This practical example envisages preventing the spindle motor from being switched on and off with high frequency when the hit rate is low.

Steps 301 through 309 are the same as steps 11 through 19 in Figure 1. In this practical example, the disk/cache memory controller 203 determines in step 310 whether the hit rate has until now been high or low when the request from the host has failed to retrieve data during step 305. Values serving as high and low reference levels should be separately preset, and the current hit rate compared with these numerical values.

If the hit rate is high, the same processing as in the practical example above will be performed. If the hit rate is low, the spindle motor 211 of the disk is switched on (if it has been switched off) in step 311. Next, in step 312, data are exchanged between the disk and the host via a cache memory 204.

Merits of the Invention

Because the present invention involves operating the spindle motor of a magnetic disk only in the case of a miss, the loss of electric power is reduced, and power consumption can be lowered.

Another advantage is that because the spindle motor always operates when the hit rate is low, it is possible to reduce the time delay caused by the switching on and off of the spindle motor.


4. Brief Description of the Drawings

Figure 1 is a flow chart showing the processing performed in accordance with a practical example of the present invention, Figure 2 is a block diagram of the hardware [for carrying out the processing shown] in Figure 1, and Figure 3 is a flow chart showing the processing performed in accordance with another practical example of the present invention.

Key

203: disk/cache memory controller, 204: cache memory, 205: encoder/decoder data separator, 206: waveform shaping circuit, 207: R/W IC, 211: spindle motor


Figure 1



Key to Figure 1

- 11 Start
- 12 Disk spindle motor off
- 13 Disk controller maintains standby mode of low power consumption until request received from host
- 14 Disk controller cancels low power consumption mode when request received from host
- 15 Request from host?
- (Below 15) Hit
- 16 Data exchange between cache memory and host
- (Above 17) Miss
- 17 Disk spindle motor on
- 18 Data exchanged between disk and host via cache memory
- 19 Disk spindle motor off

Figure 2



Key to Figure 2

- 201 Host computer
- 202 Host interface controller
- 203 Disk/cache memory controller
- 204 Cache memory
- 205 Encoder decoder data separator
- 206 Waveform shaping circuit

(Between 203 and 207) [Power save control 2]


(Above 211 to the left) Head

(Above 211 to the right) Disk

208 Voice coil motor driver
209 Voice coil motor

(To the left of 210) Power save control 1
210 Spindle motor driver
211 Spindle motor

Figure 3



Key to Figure 3

- 301 Start
- 302 Disk spindle motor off
- 303 Disk controller maintains standby mode of low power consumption until request received from host
- 304 Disk controller cancels low power consumption mode when request received from host
- 305 Request from host?
- (Below 305) Hit Data exchange between cache memory and host
- 306 (Above 310, to the right of 305) Miss
- 310 (Above 307) Hit rate?
- 307 (Above 307) High
- 308 (Above 307) Disk spindle motor on
- 309 (Above 307) Data exchanged between disk and host via cache memory
- 309 (Above 307) Disk spindle motor off
- (Above 311, to the right of 310) Low
- 311 (Above 311) Spindle motor switched on if previously off
- 312 (Above 311) Data exchanged between disk and host via cache memory